Efficient Solution of Large-Scale Electromagnetic Eigenvalue Problems using the Implicitly Restarted Arnoldi Method

Daniel White and Joseph Koning
Lawrence Livermore National Laboratory
P.O. Box 808, Livermore, CA
dwhite@llnl.gov

1. Introduction

We are interested in determining the electromagnetic fields within closed perfectly conducting cavities that may contain dielectric or magnetic materials. The vector Helmholtz equation is the appropriate partial differential equation for this problem. It is well known that the electromagnetic fields in a cavity can be decomposed into distinct modes that oscillate in time at specific resonant frequencies. These modes are referred to as eigenmodes, and the frequencies of these modes are referred to as eigenfrequencies. Our present application is the analysis of linear accelerator components. These components may have a complex geometry; hence numerical methods are required to compute the eigenmodes and eigenfrequencies of these components.

The Implicitly Restarted Arnoldi Method (IRAM) is a robust and efficient method for the numerical solution of the generalized eigenproblem $\mathbf{A}x = \lambda \mathbf{B}x$, where $\mathbf{A}$ and $\mathbf{B}$ are sparse matrices, x is an eigenvector, and λ is an eigenvalue. The IRAM is an iterative method for computing extremal eigenvalues; it is an extension of the classic Lanczos method. The mathematical details of the IRAM are too sophisticated to describe here; instead we refer the reader to [1]. A FORTRAN subroutine library that implements various versions of the IRAM is freely available, both in a serial version named ARPACK and parallel version named PARPACK.

In this paper we discretize the vector Helmholtz equation using 1st order H(curl) conforming edge elements (also known as Nedelec elements). This discretization results in a generalized eigenvalue problem which can be solved using the IRAM. The question of so-called spurious modes is discussed, and it is shown that applying a spectral transformation completely eliminates these modes, without any need for an additional constraint equation. Typically we use the IRAM to compute a small set ($n < 30$) of eigenvalues and eigenmodes for a very large systems ($N > 100{,}000$).

2. Problem Formulation

We are interested in solving the vector Helmholtz equation in a 3-dimensional inhomogeneous volume Ω,

$$\nabla \times \mu^{-1} \nabla \times E - \omega^2 \varepsilon E = 0 \text{ in } \Omega, \tag{1}$$

with boundary condition $\hat{n} \times E = 0$ on $\partial\Omega$, where E is the electric field vector, μ and ε are the tensor permeability and permittivity, and ω is the radian frequency. Employing the Galerkin procedure using 1st order edge elements denoted by W_i results in the generalized eigenvalue problem

$$\mathbf{A}e = \omega^2 \mathbf{B}e \tag{2}$$

where e is the N-dimensional vector of degrees-of-freedom and the matrices **A** and **B** are given by

$$\mathbf{A}_{ij} = \int_{\Omega} \mu^{-1} \nabla \times W_i, \nabla \times W_j \,,\ \mathbf{B}_{ij} = \int_{\Omega} \varepsilon W_i, W_j \,. \tag{3}$$

The details of computing the matrices can be found in most finite element textbooks, for example [2][3]. Using terminology from continuum mechanics, the matrix **A** is referred to as the stiffness matrix and the matrix **B** is referred to as the mass matrix. When μ and ε are real, the matrices **A** and **B** are symmetric; the matrix **A** is semi-definite and the matrix **B** is positive-definite.

3. A Comment on Spurious Modes

Equation (1) admits to two types of solutions; irrotational field solutions and solenoidal field solutions. An irrotational field is the gradient of a scalar potential function

$$E_{ir} = -\nabla \phi. \tag{4}$$

Inserting (4) into (1) we see that $\omega = 0$ for irrotational fields. Conversely, by taking the divergence of (1) we see that if $\omega \neq 0$ then the field must be solenoidal,

$$\nabla \bullet \varepsilon E_s = 0. \tag{5}$$

Since the permittivity may not be continuous, equation (5) is best understood in the weak sense: we multiply (5) by a scalar potential ϕ that is zero on the boundary, and then integrate over the domain Ω and employ the divergence theorem to yield

$$\int_{\Omega} \varepsilon E_s \bullet \nabla \phi = 0. \tag{6}$$

Equation (6) states that a solenoidal solution of (1) is orthogonal to every irrotational solution. Solutions of (1) therefore can be decomposed into irrotational ($\omega = 0$) and solenoidal ($\omega \neq 0$) solutions, with every solenoidal solution being orthogonal to every irrotational solution. An important property of the vector finite element method is that the discrete Helmholtz equation (2) has the same decomposition of solutions as the original PDE. Let $L_h \subset H_1$ be the set of standard bi-linear nodal finite element basis functions, and let the discrete scalar potential be an element of L_h. It can be shown that the finite element spaces L_h and W_h are related by $\nabla L_h \in W_h$; therefore the gradient of every nodal basis function can be written as a linear combination of edge functions [4]. Because of this there exists exactly K discrete irrotational fields that are gradients of discrete scalar potentials, where K is the

number of internal nodes in the mesh. It can also be shown that these discrete irrotational fields form the null space of the stiffness matrix A; there are exactly K solutions of (2) with $\omega = 0$. These are the so-called spurious modes; they are static solutions of (2) with non-zero divergence. They are mathematically valid solutions of (2), but they are physically uninteresting. We are interested in the solenoidal solutions of (2). As in the continuous case, the discrete solenoidal solutions of (2) are orthogonal to the discrete irrotational solutions according to the inner product $u^T \mathbf{B} v$; this is a basic property of symmetric generalized eigenvalue problems.

4. Application of ARPACK for Electromagnetic Eigenvalue Problems

In this section we discuss the application of ARPACK, which is a specific implementation of the IRAM for computing the solenoidal eigenvectors and corresponding eigenvalues of (2). ARPACK provides several different subroutines depending upon whether the mass matrix **B** is the identity or not, whether the matrices **A** and **B** are symmetric or not, whether **A** and **B** are complex valued or not, etc. The mass matrix **B** would reduce to the identity matrix if Cartesian grid finite difference discretization of the vector Helmholtz equation was used. The matrices **A** and **B** would be complex valued and non-Hermitian if μ and ε were complex valued, representing lossy dielectric and magnetic materials. In our case **A** and **B** are real symmetric matrices.

The IRAM is an iterative method for computing a small set of extremal eigenvalues. In ARPACK the user can select to compute largest algebraic eigenvalues, smallest algebraic eigenvalues, largest absolute value eigenvalues, or smallest absolute value eigenvalues. ARPACK requires $N \bullet O(m) + O(m^2)$ storage where N is the dimension of the system and m the desired number of eigenvalues to compute. The basic user-specified parameters are:

a) N, the dimension of the system
b) m, the desired number of eigenvalues
c) WHICH, a character string denoting which eigenvalues to compute
d) *ncv*, the number of Lanczos basis vectors to use
e) *tol*, the numerical tolerance used to determine convergence
f) *maxit*, the maximum number of iterations

The number of Lanczos vectors must be at least $m+1$. The optimal choice of *ncv* with respect to m is problem dependent and experimentation is required. If the eigenvalues are well separated then $ncv \approx 2 \bullet m$ is acceptable, where well separated is defined as

$$\left|\lambda_i - \lambda_j\right| > C\left|\lambda_N - \lambda_1\right| \quad \text{for all } j \neq i \text{ with } C >> \varepsilon_M . \tag{7}$$

The numerical tolerance *tol* is used in the stopping criteria $\left|\lambda_c - \lambda_t\right| \leq tol \bullet \left|\lambda_t\right|$, where λ_c is the computed eigenvalue and λ_t is the exact eigenvalue closest to λ_c. It is tempting to make *tol* a very small number, however this increases the run time of the problem. If *tol* is too small convergence may not occur. It is important to remember that the eigenvalue problem (2) is itself an approximation to physical reality, so there is little point in computing the eigenvalues of (2) exactly.

The parameter *maxit* specifies the maximum number of iterations, or restarts, of the IRAM. The iteration process begins with an initial vector v_1, which is usually chosen at random. An Arnoldi factorization is computed, and the *ncv* eigenvalues of the $ncv \times ncv$ Arnoldi matrix are computed. This represents one iteration. The salient feature of the IRAM is the ability to automatically repeat this process with improved initial vectors v_i, where each new v_i is determined by application of a polynomial in **A** to the starting vector v_1. The repeated update of the starting vector through implicit restarting is designed to enhance the components of this vector in the direction of the desired eigenvalues and damp its components in the unwanted directions. The parameters *tol* and *maxit* are not independent; a large number of iterations may be required to converge with a small tolerance. Fortunately neither *tol* nor *maxit* affect the required amount of memory, these parameters only affect the run time of the calculation.

ARPACK does not require that users actually provide the matrices **A** and **B**; instead all that is required is the action of these matrices. Specifically, the user must provide $\mathbf{A}x \to w$, $\mathbf{B}x \to z$, and $\mathbf{B}^{-1}x \to y$. For small systems the mass matrix **B** can be factored, however since we are interested in large problems we employ iterative Krylov-type methods for computing $\mathbf{B}^{-1}x \to y$. If iterative methods are used for $\mathbf{B}^{-1}x \to y$ it is essential that the residual be significantly smaller than the requested eigenvalue tolerance *tol*.

We are interested in computing only a few of the solenoidal eigenvectors and corresponding eigenvalues of (2). Based on the discussion of spurious modes we know that the distribution of the eigenvalues is as shown in Figure 1.

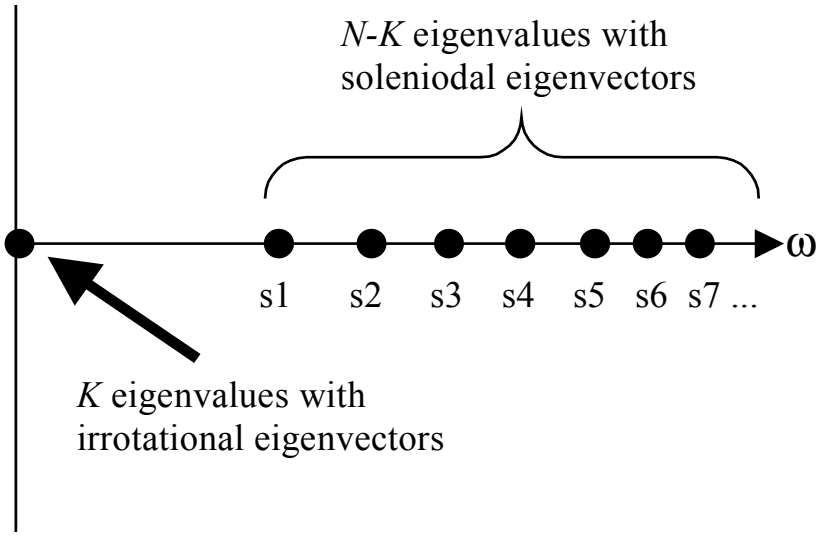

Figure 1: Eigenvalue distribution of $\mathbf{A}e = \omega^2\mathbf{B}e$

As illustrated in Figure 1, the solenoidal modes are not extremal and the IRAM cannot be directly applied to (2). Instead, a shift-and-invert spectral transformation is applied to enhance convergence to the desired part of the spectrum. If (x, λ) is an eigenpair of $(\mathbf{A}, \mathbf{B})$, and $\sigma \neq \lambda$, we form a new eigensystem

$$\mathbf{B}x = \gamma(\mathbf{A} - \sigma\mathbf{B})x$$
$$\text{with } \gamma = \frac{1}{\lambda - \sigma}. \tag{8}$$

We choose σ to be between 0 and *S1*, the smallest non-zero eigenvalue. The algebraically largest eiganvalues of the transformed system (8) correspond to the eigenvalues imediately to the right of σ in the original system (2). The zero-valued eigenvalues of the original problem are now the algebraically smallest eigenvalues of the transformed system. This is illustrated in Figure 2. The IRAM is applied to the transformed system computing the *k* algebraically largest eigenvalues and corresponding eigenvectors. The eigenvectors of (8) and (2) are identical. Once found, the eigenvalues of the original problem are computed via

$$\lambda_i = \sigma + \frac{1}{\gamma_i} \ . \tag{9}$$

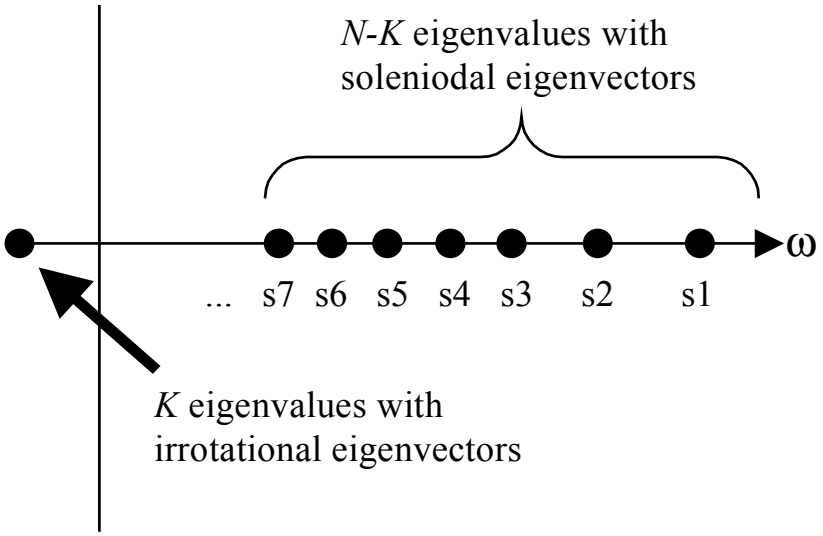

Figure 2: Spectrum of the transformed eigensystem

The shift-and-invert spectral transformation results in a new eigensystem where the desired eigenvalues are very well separated from the spurious eigenvalues. There is no need to add an additional constraint equation such as done in [5], nor is there any need to modify ARPACK. A disadvantage of this approach is that the user must have some knowledge of the problem to wisely choose σ. In addition, it is now necessary to provide the action of $(\mathbf{A} - \sigma\mathbf{B})^{-1}x \rightarrow y$. The matrix $\mathbf{A} - \sigma\mathbf{B}$ is indefinite and may be ill conditioned. However it is no more ill conditioned than the matrices that arise in standard frequency domain finite element electromagnetics where the frequency is a user-specified parameter and the right-hand side of (1) is non-zero.

5. Results

5.1 A Simple Sphere

The first problem is that of a simple homogeneous sphere. Although this seems trivial, it is in fact a difficult problem from a numerical point of view due to the numerous degenerate eigenvalues. The 20 smallest eigenvalues for a 36 cell per radius sphere are shown in Table 1, along with the exact solution. The sphere had a radius of 0.05855m and the speed of light is unity. The mesh had a total of 55296 hexahedral cells and the dimension of the eigensystem was $N = 162528$, the number of internal edges in the mesh. A sigma of $\sigma = 2000.0$ was used to perform the shift-and-invert spectral transformation.

Mode	Exact	Computed	Percent Error
TM11	2196.39	2200.44	0.184
TM11	2196.39	2200.44	0.184
TM11	2196.39	2200.44	0.184
TM21	4368.84	4382.21	0.306
TM21	4368.84	4382.21	0.306
TM21	4368.84	4384.45	0.357
TM21	4368.84	4384.45	0.357
TM21	4368.84	4384.45	0.357
TE11	5888.69	5911.07	0.380
TE11	5888.69	5911.07	0.380
TE11	5888.69	5911.07	0.380
TM31	7214.14	7248.11	0.470
TM31	7214.14	7248.21	0.470
TM31	7214.14	7248.21	0.470
TM31	7214.14	7248.21	0.470
TM31	7214.14	7252.40	0.530
TM31	7214.14	7252.40	0.530
TM31	7214.14	7252.40	0.530
TE21	9688.19	9731.82	0.450
TE21	9688.19	9731.82	0.450

Table 1: Exact vs computed eigenavlues (ω^2) for 36 cell per radius sphere

The data in Table 1 was computed using the following ARPACK input parameters:

a) $N = 162528$
b) $m = 20$
c) WHICH = largest algebraic eigenvalues
d) $ncv = 40$
e) tol = 1.0e-5
f) $maxit = 300$

Although *maxit* was set to 300, only six iterations were required for convergence of all 20 eigenvalues. The Jacobi-preconditioned conjugate residual method was used to evaluate $(\mathbf{A} - \sigma\mathbf{B})^{-1} x \rightarrow y$ with a residual tolerance of 1.0e-9. The total run time was 25 hours on a Compaq AlphaServer 8400 with a

theoretical peak performance of 880 Mflops. The run time was dominated by the cost of the conjugate residual method and not by ARPACK or by the calculation of the finite element mass and stiffness matrices.

5.2 A Linear Accelerator Induction Cell

The second problem is to compute the lowest eigenmodes of a linear accelerator induction cell. Of particular interest is the magnitude of the electric field in the accelerating gap, as this determines whether or not the particular mode will couple with the electron beam. A 33024 cell hexahedral mesh is used to model the induction cell. Part of the cell is vacuum, and another part consists of oil with a relative permittivity of $\varepsilon_r = 4.5$. The input parameters to ARPACK were the same as for the sphere problem above except that $N = 90237$ for this problem. Based on back-of-the-envelope estimations, a sigma of $\sigma = 0.001$ was chosen for the shift-and-invert spectral transformation. A total of 5 IRAM iterations were required for convergence of all 20 eigenmodes. The total run time was 15 hours on the same Compaq AlphaServer 8400. The computed eigenvalues are shown in Table 2, naturally for this problem there is no analytical solution for comparison. The column labeled ω^2 is the computed eigenvalues of (2) using unity speed of light, and the column labeled Mhz is the calibrated resonant frequencies of the induction cell.

Figure 3 shows the 1st, 5th, 13th, and 20th eigenmodes of the induction cell. The figures show electric field magnitude. Although it may be difficult to discern from the figures, the 1st and 20th modes have maximum field values in the accelerating gap and hence will couple strongly with the electron beam, whereas the 5th and 13th modes are examples of modes that will not couple strongly with the beam.

ω^2	Mhz
.00495	132.200
.01158	202.166
.01158	202.166
.01721	246.471
.02222	280.015
.02222	280.015
.02930	321.560
.02930	321.560
.03585	355.683
.03585	355.683
.05206	428.607
.05206	428.607
.05466	439.197
.05917	456.955
.05921	457.109
.05929	457.417
.05929	457.417
.07053	498.881
.07053	498.881
.08034	532.453

Table 2 : The 20 lowest eigenvalues and resonant frequencies for the induction cell

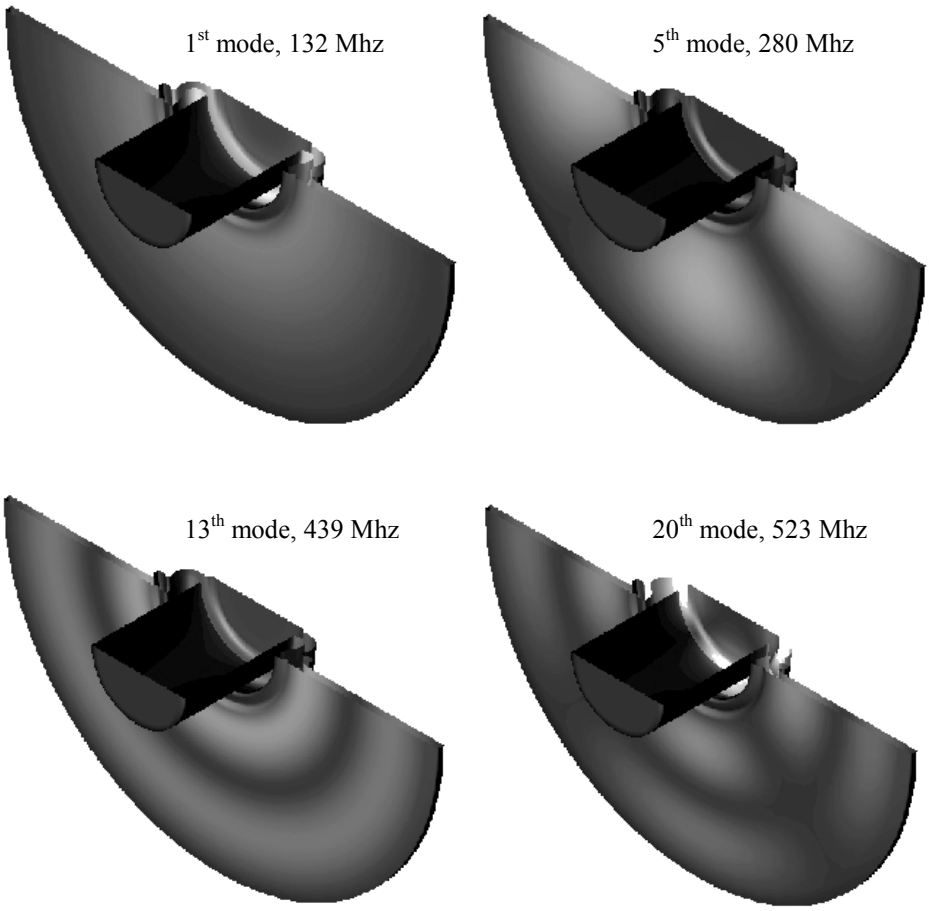

Figure 3: Selected eigenmodes for the linear accelerator induction cell. The geometry is clipped so that the internal structure of the modes can be seen.

6. References

[1] R. B. Lehoucq, D.C Sorensen, and C. Yang, ARPACK Users Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods, SIAM, 1998.

[2] J. L. Volakis, A. Chatterjee, L. C. Kempel, Finite Element Method for Electromagnetics: Antennas, Microwave Circuits, and Scattering Applications, IEEE Press, 1998.

[3] A. Bossavit, Computational Electromagentics: Variational Formulations, Complementarity, Edge Elements, Academic Press, 1998.

[4] J. C. Nedelec, "Mixed Finite Elements in R3", Numer. Math., vol. 35, pp. 315-341, 1980.

[5] S. G. Perepelitsa, R. Dyczij-Edinger, J. F. Lee, "Finite Element Analysis of Arbitrarily Shaped Cavity Resonators Using H1(curl) Elements," IEEE Trans. Mag., vol. 33, no. 2, pp. 1776-1779.